



GridGain In-Memory Data Fabric: Powering Bitcoins and Blockchain With In-Memory Computing

ERIC KARPMAN
INDEPENDENT CONSULTANT

www.gridgain.com



[#gridgain](https://twitter.com/gridgain)

Gartner 2014
CoolVendor



Current Issues with Financial Ecosystem

- Lack of trust
- Legacy systems and processes
- Too many intermediaries
- Slow and complex post-trade functions
- Lots of reconciliation
- Regulatory compliance
- Financial fraud

What is Bitcoin?

- It's a crypto-currency without any central bank authority that uses a unique mechanism of mining to introduce new currency into circulation by winning a computational race to collect a prize
- 70% of bitcoin mining is done in China (Bloomberg)
- Started as electronic micro-payments idea
- Bitcoins are also exchanged for traditional currency on virtual exchanges
- Bitcoins has surged more than 4,000% since 2011 (Bloomberg)
- Bitcoins are stored in electronic wallets
- Bitcoins are stored using Blockchain technology

What is blockchain and why now?

Sequential, permanent, encrypted, decentralized ledger of transactions, verified and confirmed by random computers anywhere within the network and distributed to any subscriber.

“75% expect widespread adoption of blockchain within the next three to six years”

“Two-thirds expect the introduction of blockchain to produce substantial savings, possibly up to 25%”

“48% argue that blockchain technology could also help the financial industry cope with the risk of system failures and market disruption”

“87% of market participants believe blockchain technology will completely change the settlement process for securities”

– Deutsche Bank and FT Remark Survey of 200 financial industry participants, November 2016

- ✓ The Internet of Trust
- ✓ Cloud
- ✓ Tendency to decentralization and peer-to-peer

Current Problems with Bitcoins and Blockchains

- User-unfriendly software
- Hacking
- System failure
- High volatility
- Limited acceptance
- Password-based authentication
- Passwords are failing
- Governmental scrutiny
- Lack of legal framework
- Too much mining power in one country
- Trust comes with time

Examples of Blockchains Used in Banking

- Netting, Matching and Clearing
- Collateral Management
- Payments Processing
- Trade Finance
- Same Day Settlement
- Securities Servicing and Processing
- Transfer Agency and Registrar Functions
- Equilibrium and decentralization of players
- Private vs Public Blockchains: registration vs. free entrance

Processing Letters of Credit

- Importer files for LC
- Importer bank approves LC
- Importer bank sends LC to exporter bank
- Exporter bank approves LC
- Exporter ships goods and sends invoice to importer
- Importer sends invoice to importer bank
- Importer bank reviews the invoice and sends to exporter bank
- Exporter bank reviews and accepts the invoice
- Exporter and importer banks settle the transaction (with many settlement steps)

Distributed Post-Trade Functions

Trade enrichment -> regulatory reporting -> prime
brokerage -> allocation -> matching -> affirmation ->
confirmation -> netting -> collateral management ->
compression -> default management -> novation ->
custody -> settlement -> asset servicing

Evolution of In-Memory Grid Computing

- Move from Disk to 100% In-Memory (RAM)
- Leverage Clustered Memory and Parallel Distributed Processing
- Results: 100x Faster, 10x ROI Improvement
- Making “Big Data” Fast

“In-memory will have an industry impact comparable to web and cloud.”

“RAM is the new disk, and disk is the new tape.”

Gartner

In-Memory Computing Market:

- \$10B in 2019
- CAGR 22%

Gartner

Financial Customer Use Cases

Data Velocity, Data Volume, Data Consistency, Real-Time Performance and Analysis

- **Trading Platforms**
Order Management and Execution Management Systems, algorithmic trading, high volume transactions, ultra low latencies.
- **Risk Management**
Modeling, financial engineering, pricing, hedging, what-if analysis.
- **Financial Analytics**
Real time analysis of trading positions, trending, market data analysis, sentiment analysis, complex event processing, hedging, transaction cost analysis, time series, volatility analysis, Monte Carlo simulations, Black-Scholes, derivatives pricing.
- **Big Data Analytics**
Customer and counter party 360 view, master data management, securities masters, reference data, real-time analysis of P&L, up-to-the-second operational BI.
- **Compliance and Monitoring**
Fraud, AML, KYC, market manipulation and abuse, pre and post trade compliance modeling.
- **Financial SaaS Platforms**
High performance next-generation architectures for Software as a Service Application vendors.



THOMSON REUTERS



Jefferies



APOLLO

Julius Bär



CAMBRIDGE



ASSOCIATES




Use Case:



Largest bank in Russia and Eastern Europe, and the third largest in Europe

- Sberbank Requirements
 - Migrate to data grid architecture
 - Minimize dependency on Oracle
 - Move to open source
- Why GridGain Won
 - Best performance
 - 10+ competitors evaluated
 - Demonstrated best
 - Fault tolerance & scalability
 - ANSI-99 SQL Support
 - Transactional consistency
 - Strict SLAs
 - Less than 5 min cluster restart
 - Fully Operational from disk & memory



**1 Billion
Transactions per
Second**

**10 Dell R610 blades = \$25K
1 TB Memory**

From RBC article – January, 2016



Herman Gref
CEO & Chairman, Sberbank


“The new Sberbank IT plan is to create a platform that enables the bank to introduce new products in hours, not weeks. The platform will have virtually unlimited performance and very high reliability. It will be much cheaper and will significantly reduce human interaction during customer transactions. The system will use machine-learning, flexible pricing, and artificial intelligence,” said German Gref, head of Sberbank.

“The new system will use technology from GridGain, which won the tender from Oracle, IBM and others, and turned out to deliver an order of magnitude higher performance than those of the largest companies,” he added.

Use Case:

Smart Metering and Utilities – delivers a comprehensive IOT platform

- SilverSpring Requirements
 - Migrate to in-memory processing
 - Add scalability & elasticity
 - Use open source technologies
- Why GridGain Won
 - Strong compute capabilities
 - Co-located in-memory processing
 - Demonstrated best
 - On-demand elasticity & scalability
 - ANSI-99 SQL Support
 - Transactional consistency

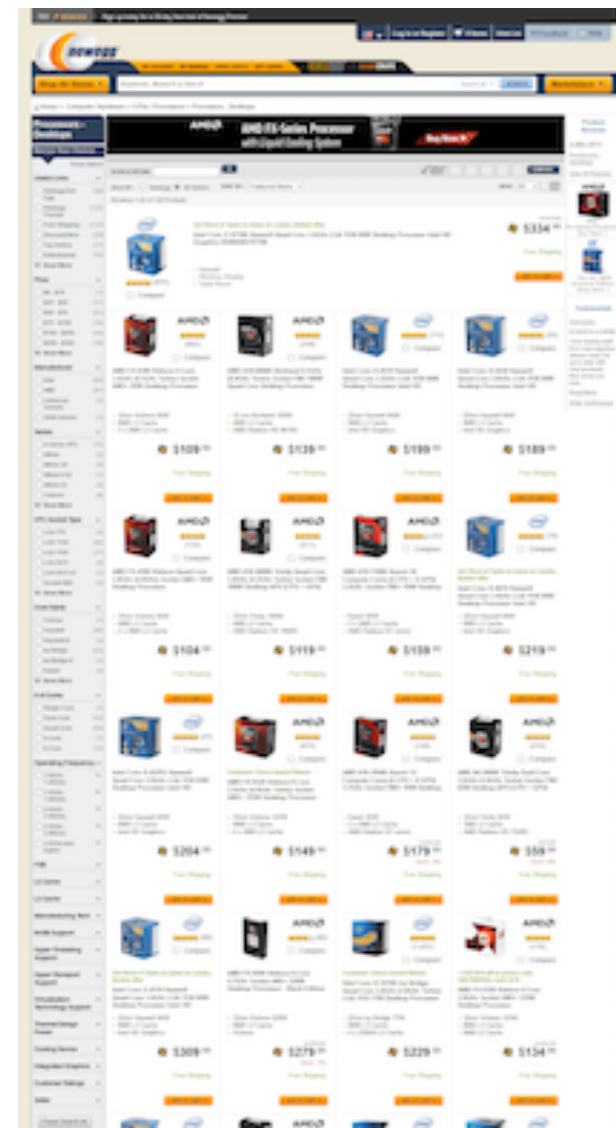


Use Case:

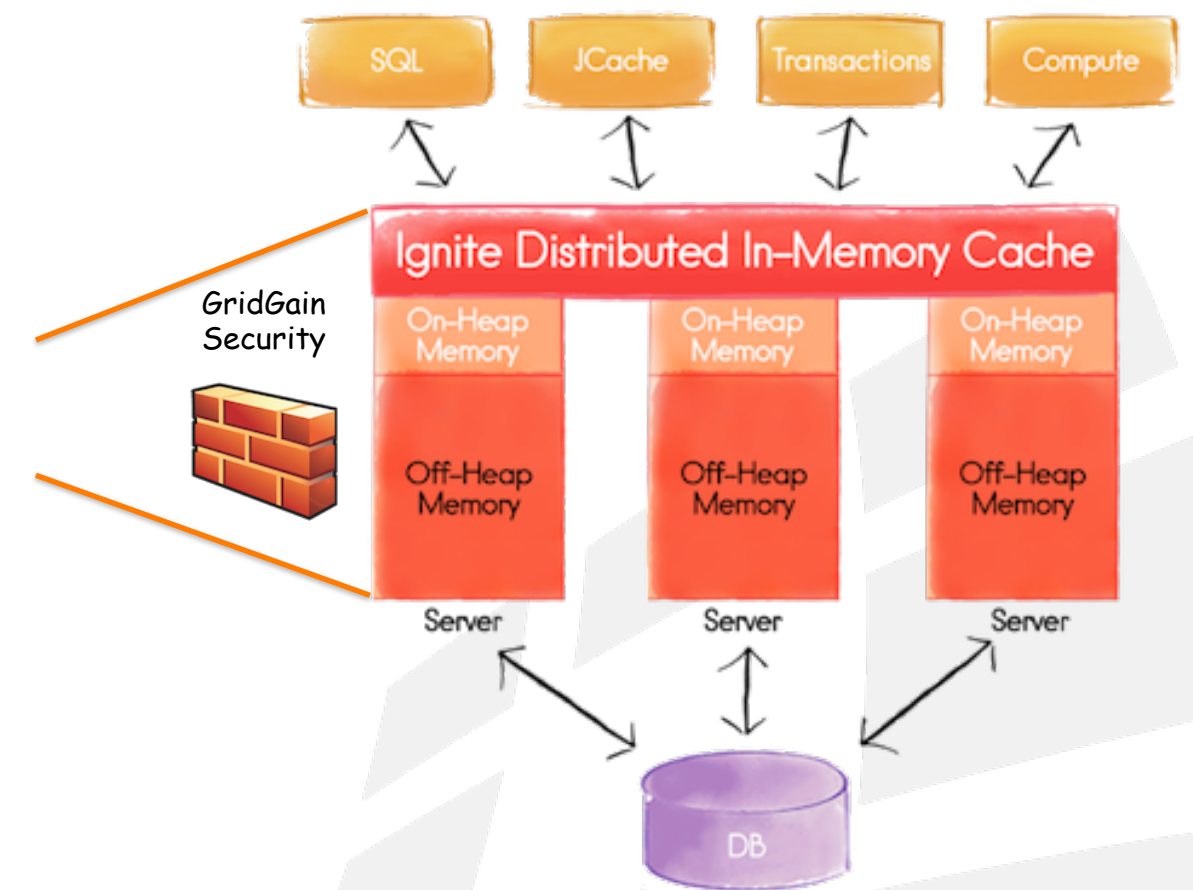


Online consumer electronics store – 11th largest retailer in US

- Newegg Requirements
 - Oracle Coherence replacement
 - Smart in-memory analytics
 - Use open source technologies
- Why GridGain Won
 - Best performance
 - 2x better than incumbent
 - Demonstrated best
 - ANSI-99 SQL Support
 - Resiliency & fault tolerance
 - Transactional consistency




NewEgg Ecommerce Site




Why Now?

Data Growth and Internet Scale Driving Demand



8 zettabytes in 2015 growing to 35 in 2020

Declining DRAM Cost Driving Attractive Economics




Cost drops 30% every 12 months

GridGain Enterprise and Open Source Strategy




- GridGain Enterprise Edition is based on Apache Ignite
- Open source is intended to provide an easy entry point for learning, testing and non-critical use
- Enterprise Edition customers benefit from many exclusive enterprise-class features along with support and indemnification

What is an In-Memory Data Fabric?

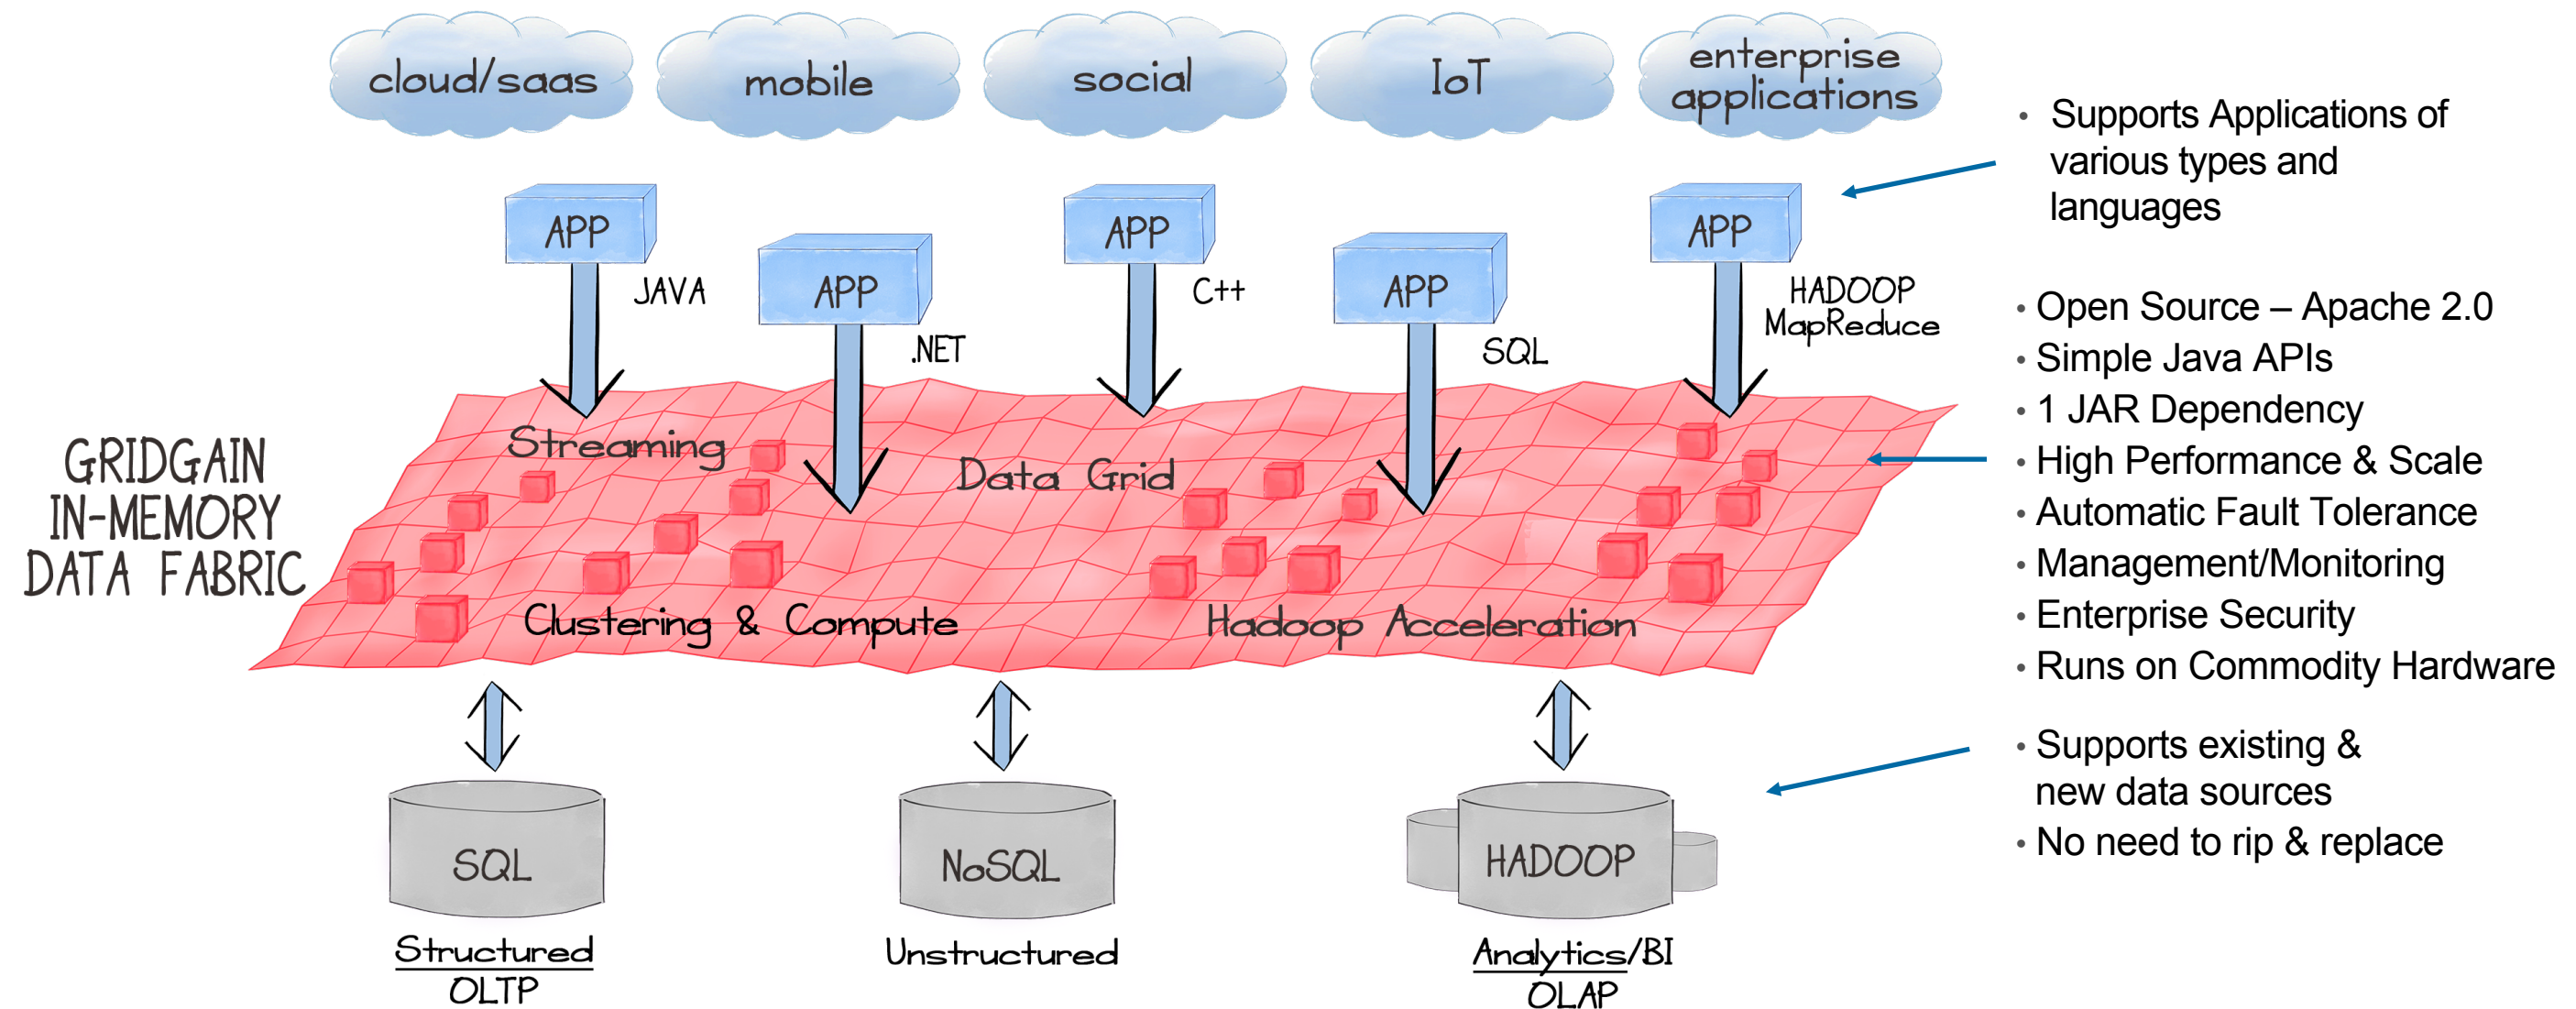


High-performance distributed in-memory platform for computing and transacting on large-scale data sets in near real-time.

In-Memory Data Fabric Use Cases



Flexibility and Enterprise Breadth of In-Memory Data Fabric



In-Memory Data Fabric: Main Benefits

- Performance
 - High Throughput
 - Low Latency
 - Load Balancing
 - Caching
 - In-Memory Indexing
 - Eliminate Java Garbage Collection Pauses
- Scalability
 - Add Cluster Members (cores)
 - Unlimited Vertical Scale
 - Add Memory (RAM)
- High Availability
 - Data Backups
 - Datacenter Replication
 - Automatic Failover
 - Persistence
 - Fault Tolerance
 - Fast Recovery



- Transactions
 - Fully ACID Compliant
 - Optimistic & Pessimistic
 - Data Streaming
- Persistence
 - SQL, NoSQL, Hadoop
 - Tiered Memory: On-Heap ->
- Security
 - Auth & Auth
 - Encryption
 - Tracing & Auditing



GridGain's In-Memory Data Fabric Enterprise Edition

GridGain Enterprise Subscriptions include:

- > Right to use GridGain Enterprise Edition
- > Bug fixes, patches, updates and upgrades
- > 9x5 or 24x7 Support
- > Training and Consulting Services from GridGain

Features	Apache Ignite	GridGain Enterprise
<i>In-Memory Data Grid</i>	✓	✓
<i>In-Memory Compute Grid</i>	✓	✓
<i>In-Memory Service Grid</i>	✓	✓
<i>In-Memory Streaming</i>	✓	✓
<i>In-Memory Hadoop Acceleration</i>	✓	✓
<i>Distributed In-Memory File System</i>	✓	✓
<i>Advanced Clustering</i>	✓	✓
<i>Distributed Messaging</i>	✓	✓
<i>Distributed Events</i>	✓	✓
<i>Distributed Data Structures</i>	✓	✓
<i>Portable Binary Objects</i>	✓	✓
<i>Management & Monitoring GUI</i>		✓
<i>Enterprise-Grade Security</i>		✓
<i>Network Segmentation Protection</i>		✓
<i>Recoverable Local Store</i>		✓
<i>Rolling Production Updates</i>		✓
<i>Data Center Replication</i>		✓
<i>Integration with Oracle GoldenGate</i>		✓
<i>Basic Support (9x5)</i>	✓	✓
<i>Enterprise Support (9x5 and 24x7)</i>		✓
<i>Security Updates</i>		✓
<i>Maintenance Releases & Patches</i>		✓

Free
w/ optional Paid Support

Annual License
Subscription



THANK YOU